Security Infrastructure for Grid-Enabled Biomedical Services

Anthony Stell, Micha Bayer, Jos Koetsier, Richard Sinnott
National e-Science Centre
University of Glasgow & University of Edinburgh
ajstell@dcs.gla.ac.uk
michab@dcs.gla.ac.uk
jos@nesc.ac.uk
ros@dcs.gla.ac.uk

Abstract

The potential benefits and advantages that Grid Computing can bring to society in different fields of human endeavour are many and varied. These benefits however, can only be usefully realised if they can be implemented securely. Currently, there are many standards in the Grid community for enforcing security, particularly authorization, and no single specification has been adopted as the definitive solution. This leads to the situation where the standards and specifications that are available must be rigorously tested, applied and compared before being used in a production context. This paper outlines applied security measures undertaken in the BRIDGES project (Biomedical Research Informatics Delivered through Grid-Enabled Services) that exemplify these issues. A discussion of the security problems, inherent in both the computational and the data strands of this project, is presented along with the solutions implemented at NeSC (National e-Science Centre).

1. Introduction

The BRIDGES project [1] conducted at the National e-Science Centre has a focus on delivering a Grid infrastructure offering secure access to and usage of highly distributed, evolving biomedical data sets for the Wellcome Trust funded Cardiovascular Functional Genomics (CFG) consortia [2]. The CFG scientists wish to:

- seamlessly access public genomic data resources;
- securely share data sets with one another;
- get simplified, secure access to high performance computing resources, for example in order to run Basic Local Alignment Search Tool (BLAST) applications.

The work in BRIDGES therefore falls broadly into two categories: computation and data retrieval. The computation strand concerns the development of a GridBLAST service that runs over several machines but is presented to the user as a single process. The data retrieval strand follows the Grid paradigm by retrieving information from several distributed databases, federating the data and returning it to the user as if it were a single resource.

The CFG Virtual Organisation (VO) is outlined in figure 1.

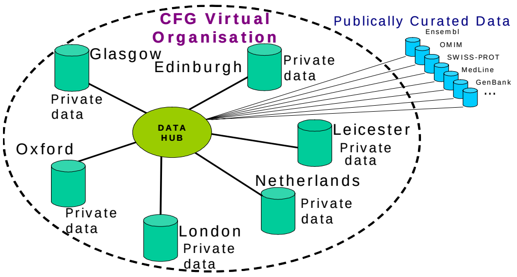

Figure 1: CFG Virtual Organisation

For data retrieval, access to public genomic data sets such as Ensembl (rat, mouse, human databases) [3] and MGI [4] has been achieved through the use of IBM's Information Integrator and the Grid community's OGSA-DAI [19] technology. The numerous other databases the scientists need access to could only be solved by the establishment of a local warehouse (including OMIM, HUGO, RGD GO) [5-8], as programmatic access to these databases is unavailable. The Data Hub (in figure 1) is based upon DB2 and has user friendly front end client tools available via the BRIDGES portal (GeneVista, MagnaVista [1]).

These tools provide both a customisable front end to the numerous data sets and the interface to the security infrastructure. Typically, the CFG scientists input gene names which result in the return and display of all of the data sets associated with those genes from all of the aforementioned databases. We note that the data returned is configurable, i.e. users can identify which data sets should be returned, which databases should be queried etc.

2. Security Challenges

To efficiently utilise Grid technology, services must be secured using effective yet flexible means. The overall structure required is similar to the security used for most service providers: it consists of a user-access portal, an authentication mechanism, an authorization mechanism and has a provision for logging and activity management.

Common to both of the strands, in addition to authentication for using the services, is the need for a dynamic policy that has the ability to evolve with time but in such a manner as not to disrupt any current operation (this is a central tenet of Grid Computing). Issues of application scalability and delegation of security control between VO and the local resource administrator also need to be addressed.

Additionally, there are security demands imposed by distributed job submission and data integration that differ subtly in their nature and their necessary solutions.

2.1 Compute security

The security demands of job submission require that only users that are authorised may use specific resources. This should be controlled by the service provider and must be an automated process - otherwise it would not be usefully exploiting the advantages of Grid technology.

The main abuses of computational resources that could occur include the oversubscription of service resources, which would disrupt the nodes being used, and the provision of access to users who are unknown to the institution.

Control of this requires identification of every user on the system and enforceable authorization policies that allow strict control of resource allocation by the service provider only.

2.2 Data security

The problems inherent in enforcing data security are more difficult to overcome using current technology. The ideal scenario is that the data within the database can be secured based on the granularity of the database schema itself (for instance, per-table or per-row). However, this requires retrieving data, or meta-data, from the database using SQL first. So a “Catch-22” is arrived at: the data must be retrieved first to find out whether the user is allowed to retrieve the data.

A possible solution to this problem is the use of per-parameter authorization (as opposed to per-method authorization). This would require a change to the SAML specification that technologies such as PERMIS [10] are built upon, and a GGF request to amend this specification has recently been authorised [16]. The extra information that could be passed along with the authorization request would allow the user credentials to be matched with the authorization control enforced by the resource itself, or possibly by a trusted third party.

The interim higher-level solution therefore, is that the different databases, from which the information is drawn, can have authorisation restrictions imposed. So the basic policy allows those within the VO (in this case the CFG consortium) access to all federated databases whilst those outside that VO only have limited access to these databases depending on their local restrictions. This lacks the fine granularity of the ideal data security policy, but at least provides a form of authorization at the VO level.

3. Proposed Security Solutions

A standard security infrastructure is needed to initially ring-fence the services provided by BRIDGES.

A web portal must be provided that restricts access to the services and enforces one route that clients must go through to access the services beyond. Then once clients have been authenticated, access control is enforced, based on their credentials, using a form of Role-Based Access Control (RBAC).

Finally, a logging mechanism is used that records activity by users that are accessing the BRIDGES services to allow auditing and enforce accountability.

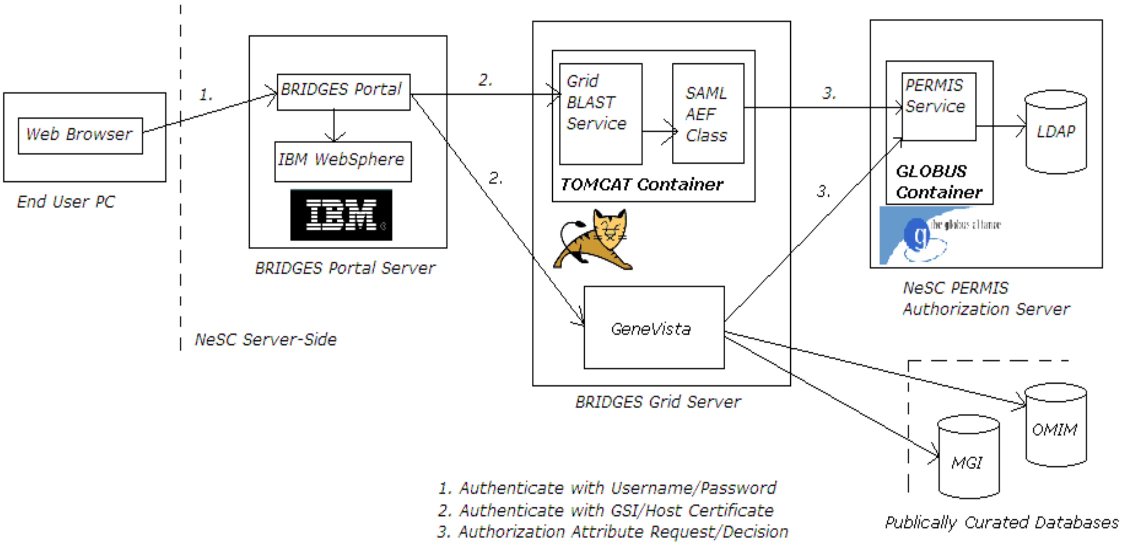

Figure 2: BRIDGES Security Infrastructure

3.1 Comparison with existing solutions

Beyond the standard security mechanisms there are alternative solutions that exist for enforcing Grid security. Username/password combinations exist as an adequate way of providing authentication but that is where their usefulness ends. A more sophisticated method of attaching privileges to user identity is needed to enforce access control. The authorization technology that is in the most widespread use just now is the Grid Security Infrastructure (GSI) method of constructing a list of Distinguished Names (DNs) against local usernames. This acts simply as a lookup table, which the service consults when it is asked for an authorization decision. Again, this does not provide enough flexibility to be used in the context of Grid technology.

Two criteria that are most important when dealing with access control are those of scalability and delegation of trust. There are a number of solutions that attempt to address these in slightly different ways. CAS [17] looks at the issue of trust delegation between the VO and the local resource, whilst VOMS [18] attempts to manage scalability by automatically generating a set of inter-linked gridmap files.

The solution used here is the PERMIS technology, which approaches the scalability problem by using an LDAP repository of policies and looks at trust delegation by structuring those policies with the use of hierarchical XML statements.

3.2 Compute Security Implementation

In BRIDGES, the PERMIS authorization policy controls what resources the job can be run on for that particular user. Currently the resources available for use include the NGS [13], ScotGrid [14] and the local NeSC Condor pool. It is not mandatory for end users to have a UK e-Science certificate - access to resources such as the NGS is permitted through the use of a single host certificate on the grid server. The security infrastructure automatically identifies the users and submits their jobs to the appropriate resources (as defined in the XML policy). Currently three policies are supported:

- If a user has invalid or no explicit privileges to present then they can run their jobs on the local Condor pool at NeSC.
- If a user has a ScotGrid account at NeSC, then they can use this resource.
- If a user has been vetted by the local NeSC management of these services then they can use the Grid resources provided by the NGS.

As long as the specific criteria for each resource are met then a user can have any combination of these three resources. This is defined in the XML policy, stored in the LDAP repository.

3.3 Data Security Implementation

PERMIS policies have been defined and implemented restricting access to certain databases, offered via the GeneVista application, to certain users. This is achieved through extensions to the application layer of GeneVista, which indirectly supports queries of the PERMIS-based policies.

The extensions include a boolean parameter to define whether a user is authorised to access a specific database or not. Depending on the user credentials presented, the parameter will be switched and the database will be added to the list of those from which data must be retrieved.

3.4 Infrastructure and Status

The overall secure setup is shown in figure 2. The three servers involved incorporate a portal server, a grid server and an authorization server. The portal server, implemented using IBM Websphere [9], provides the necessary authentication through the familiar mechanism of a username/password combination.

The grid server hosts the BLAST service, implemented using version 3.0 of the Globus Toolkit [12]. The GeneVista application is run from a portal server. Both of these services call out to the authorization server for access control decisions.

The authorization server hosts the PERMIS service, which is deployed in a Globus container (v3.3). The two call-outs (BLAST and GeneVista) use the same policy but the targets are defined separately within this policy. Extra policies can be defined and added to the LDAP repository, if needed.

The federated data and processed BLAST service return their results to the portal server which allows the client to view the results. The compute resources are made available at the resource-level by providing users with accounts on ScotGrid and using a server certificate to access the NGS. User management is such that users will not be given access to these resources in the XML policy if they do not have the necessary privileges for the underlying resource.

The data authorization is implemented in a simpler fashion. The GeneVista application federates the data, gathered from publicly curated databases, and sends this information back to the portal, which renders it to the client. Authorization requests, asking for a decision about this user on each database, are sent to the authorization server before the data is retrieved and, depending on the results of this, a limited table of data is rendered by GeneVista. The back-end infrastructure for GeneVista has been set up to use test users as this is a proof-of-concept scenario.

4. Conclusions and Future Work

This paper gives an outline of how advanced Grid security infrastructures are applied and supported within the BRIDGES project. Future work includes the migration of this code structure to make use of the latest release of the Globus Toolkit (version 4.0). The structures used will also be applicable to other projects doing similar work of securing access to data and computation. The MRC-funded VOTES project (Virtual Organisations for Trials and Epidemiological Studies) [15] will make direct use of the solutions developed here to provide flexible, effective security to distributed data concerning clinical trials.

The GGF Authorization and Authentication Working Group is working on the development of a SAML standard that will allow per-parameter authorization. This will be useful in the context of exercising authorization controls over the actual content of the databases, as opposed to the databases themselves. Once this specification is complete then the development of security on GeneVista can take on a more sophisticated role.

5. References

[1] BRIDGES – http://www.brc.dcs.gla.ac.uk/projects/bridges
[2] CFG – http://www.brc.dcs.gla.ac.uk/projects/cfg
[3] European Bioinformatics Institute – http://www.ebi.ac.uk
[4] Mouse Genome Informatics – http://www.informatics.jax.org
[5] NCBI Online Mendelian Inheritance in Man – http://www.ncbi.nlm.nih.gov/OMIM
[6] Human Genome Organisation – http://www.gene.ucl.ac.uk/hugo
[7] Rat Genome Database – http://rgd.mcw.edu
[8] Gene Ontology – http://www.geneontology.org
[9] IBM Websphere – http://www-306.ibm.com/software/websphere
[10] PERMIS – http://sec.isi.salford.ac.uk
[11] OASIS Assertions and Protocols for the OASIS Security Assertion Markup Language (SAML) v1.1, 2 Sep 2003 – http://www.oasis-open.org/committees/security
[12] Globus Toolkit – http://www.globus.org
[13] National Grid Service – http://www.ngs.org
[14] ScotGrid – http://www.scotgrid.ac.uk
[15] Virtual Organisations for Trials and Epidemiological Studies – http://www.nesc.ac.uk/hub/projects/votes
[16] GGF Ogsa-Authz Working Group – https://forge.gridforum.org/projects/authz-wg
[17] L. Pearlman et al., A Community Authorization Service for group collaboration, proceedings of IEEE 3rd International Workshop on Policies for distributed systems and networks, 2002
[18] VOMS Architecture, European Datagrid Authz Working Group, 5 Sep 2002
[19] OGSA-DAI – http://www.ogsadai.org